

The following paper entitled “Highly Scalable Matching Pursuit Decomposition Algorithm” has been accepted for publication in The 7th International Workshop on Structural Health Monitoring (IWSHM).

Christensen, D., S. Das, and A. Srivastava. “Highly Scalable Matching Pursuit Decomposition Algorithm.” International Workshop for Structural Health Monitoring (2009).

Highly Scalable Matching Pursuit Signal Decomposition Algorithm

D. Christensen¹ S. Das² A. N. Srivastava³

ABSTRACT

In this research, we propose a variant of the classical Matching Pursuit Decomposition (MPD) algorithm with significantly improved scalability and computational performance. MPD is a powerful iterative algorithm that decomposes a signal into linear combinations of its dictionary elements or “atoms”. A best fit atom from an arbitrarily defined dictionary is determined through cross-correlation. The selected atom is subtracted from the signal and this procedure is repeated on the residual in the subsequent iterations until a stopping criteria is met.

A sufficiently large dictionary is required for an accurate reconstruction; this in return increases the computational burden of the algorithm, thus limiting its applicability and level of adoption. Our main contribution lies in improving the computational efficiency of the algorithm to allow faster decomposition while maintaining a similar level of accuracy. The Correlation Thresholding and Multiple Atom Extractions techniques are proposed to decrease the computational burden of the algorithm. Correlation thresholds prune insignificant atoms from the dictionary. The ability to extract multiple atoms within a single iteration enhances the effectiveness and efficiency of each iteration. The proposed algorithm, entitled MPD++, is demonstrated using real world data.

INTRODUCTION

The MPD algorithm has been extensively used to decompose signals for a wide variety of applications. A special application is Structural Health Monitoring (SHM),

¹D. Christensen is currently associated with Department of Aerospace Engineering, Iowa State University, IA 50011. (email: aerodc@iastate.edu)

²S. Das is currently associated with University Affiliated Research Center, UC Santa Cruz, NASA Ames Research Park, Moffett Field, CA 94035. (email: Santanu.Das-1@nasa.gov)

³A. N. Srivastava is the Principal Investigator of the Integrated Vehicle Health Management Project at NASA. He is in the Intelligent Systems Division at NASA Ames Research Center, Moffett Field, CA, 94035 (email: Ashok.N.Srivastava@nasa.gov)

the primary motivation behind this paper [1]. In terms of SHM, primary tasks focus on the detection, localization, classification, and prediction of remaining useful life (RUL) of structures and materials. Numerous existing techniques, such as x-ray scattering, Raman spectroscopy, nuclear magnetic resonance, gamma ray, eddy currents, thermography, mass spectrometry, elastic wave propagation, and spectrum of surface waves are applicable for SHM. All these techniques may have different physical mechanisms governing their behavior; however, the same principle of analyzing sensor data to detect changes in material properties or boundary conditions applies.

In a faulty (or damaged) state, changes in material properties or structural integrity cause a deviance from the baseline, healthy-state system response. As a result, extracting the characteristic signatures regarding the quality of health of the structure or material may reveal valuable information for the prognostics and diagnostics of the system. Feature extraction detects the primary components of the signal while neglecting much of the inherent noise. However, the necessity of precision in the decomposition complicates the feature extraction algorithm. Failure to avoid mediocre or inaccurate decompositions may result in the false classification of the system's health.

Numerous feature extraction algorithms exist, including the discrete Fourier Transform, wavelet transform, MPD, and maximum entropy method. Each of these techniques have their strengths and weaknesses and some are better adopted for specific applications. It is important to note that the current research is not intended to make any comparison between the above mentioned techniques, rather the key issue that has been addressed is the "time-complexity" limitation of MPD which has drawn considerable interest in recent times.

Although MPD tends to be computationally inefficient for many applications, the dictionary elements may have fewer restrictions compared to other methods, especially regarding element orthogonality. This research is focused to improve the computational efficiency and scalability of the MPD algorithm so that it may be applied more effectively. Two unique techniques are proposed to accomplish this goal.

CLASSICAL MPD THEORY

The MPD algorithm is an iterative, nonlinear algorithm that decomposes a signal into a linear expansion of function segments or "atoms" that belong to a redundant dictionary [2]. A composite dictionary is comprised of arbitrarily defined atoms. During each iteration, the cross-correlation is calculated for every atom in the dictionary to quantify the degree of similarity between atoms and the signal. Each atom is normalized to unit energy. As a result, the amplitude of a match is merely the cross-correlation value of the matched atom with the specified portion of the signal. The dictionary's best match, the atom with the largest cross-correlation value, is extracted from the signal and stored for later reconstruction or classification. MPD is an greedy algorithm that extracts the largest amount of energy possible per iteration. Greedy algorithms conduct each iteration independently with no knowledge of previous iterations and with no concern for future iterations. The numerical procedure for the MPD algorithm is shown in the following pseudocode.

Algorithm 1 Matching Pursuit Decomposition

- 1: **Build Dictionary:** $D = \{d_{\gamma_1}, d_{\gamma_2} \dots, d_{\gamma_j} \dots, d_{\gamma_n}\}$, where $d_{\gamma_n}(t) = \frac{1}{\sqrt{\alpha_n}} d \left(\frac{t - \beta_n}{\alpha_n} \right) e^{j2\pi\kappa_n t}$
 - 2: **Initialize** $K_{stop}, \delta_{stop}, k = 0, R_x^0[n] = x[n], E_x^0 = \|R_x^0\|_2$
- MPD routine**
- 3: **while** $k < K_{stop}$ or $E_x^k > \delta_{stop}$ **do**
 - 4: $a_{\gamma_j}^k = \langle R_x^k[n], d_{\gamma_j}^k \rangle$
 - 5: **Select** dictionary element whose time correlation with the $R_x^k[n]$ is maximum
 - 6: $R_x^k[n] = R_x^{k-1}[n] - a_{\gamma_j}^k d_{\gamma_j}^k[n]$
 - 7: $k = k + 1, E_x^k = \|R_x^k\|_2$
 - 8: **end while**
-

The residual after k iterations is represented by $R_x^k[n]$, where the signal is notated by $x[n]$. In Algorithm 1 (step 2), $R_x^0[n]$ represents the initialized state prior to the first iteration. The best fit atom from dictionary D is determined through cross-correlation with the residual as shown in Eqn. 1. The selected atom is subtracted from the signal (step 6) and the procedure is repeated on the residual in the subsequent iteration. The i^{th} matched atom is represented by d_{γ_i} and its corresponding cross-correlation with a_{γ_i} , where the parameters for a particular atom are notated by γ .

$$a_{\gamma_i} = \operatorname{argmax}_{d_{\gamma_i} \in D} \left| \langle R_x^k[n], d_{\gamma_i} \rangle \right| \quad (1)$$

The algorithm is repeated until a stopping criteria is reached: a specified number of matches are removed from the signal or a designated amount of energy is extracted from the signal. These criterions establish a definition for the completion of the decomposition and their values are typically determined by the nature of the application. Once decomposed, the modeled signal \hat{x} may be reconstructed using Equation 2.

$$\hat{x} = \sum_{i=1}^k a_{\gamma_i} d_{\gamma_i}. \quad (2)$$

PROPOSED ALGORITHM: MPD++

Related research has been conducted to improve the computational performance of MPD including: coarse-fine grids, interpolation, and particle filters [1] [3]. Several other research efforts in this area may be found in the following literatures: [4] [5] [6] [7]. A common approach involves the implementation of an intelligent, self-adaptive, or reduced-in-size dictionary to assuage the complications caused by over-complete dictionaries. However, none of the above techniques implement a feedback loop between the iterations and the dictionary. This concept has been advanced to dynamically reduce the size of the dictionary based on the results of the previous iteration. Our second effort deals with the extraction of multiple atoms per iteration. To the best of the authors' knowledge, no prior research discussing multiple atom extractions has been published.

Sparsifier Dictionary with Correlation Thresholding

The MPD algorithm's bottleneck lies in the computation of the atoms' cross-correlation values. The operations required to compute the cross-correlations for L atoms with a signal of dimension M have a high complexity of $O(LMN)$, where N is the length of the reference signal. The complexity may be reduced by decreasing the size of the dictionary. While a sufficiently large dictionary is required for accurate decomposition, a majority of the atoms may not be used throughout the decomposition process. The removal of these unused atoms from the dictionary potentially offers significant increases in performance with no cost in accuracy.

Correlation Thresholding (CT) begins by calculating the correlation values for every atom in the dictionary through a standard MPD iteration. Next, the correlation ratio for each atom is calculated. The correlation ratio, CR, for a particular atom is defined as that atom's maximum correlation value divided by the largest correlation in the dictionary (Eqn. 3).

$$CR_n = \frac{\langle R_x^k[n], d_{\gamma_i} \rangle}{\operatorname{argmax}_{d_{\gamma_i} \in D} |\langle R_x^k[n], d_{\gamma_i} \rangle|} \quad (3)$$

CT applies a threshold to the correlation ratio for each atom in the dictionary. Atoms with correlation ratios above the threshold form the reduced dictionary while the rest are rejected. This reduced dictionary is used in the subsequent iterations instead of the full, base dictionary. With fewer atoms to analyze, the algorithm requires less time per iteration. After a re pruning criteria is met, the reduced dictionary will be revoked and regenerated from the base dictionary. The performance of the CT modification is governed by three thresholds:

- **Reduction Threshold:** This threshold specifies the minimum correlation ratio admitted into the reduced dictionary. The following two thresholds are determined relative to but operate independently of the reduction threshold. After the construction of the reduced dictionary, the reduction threshold plays no role in the algorithm until the next re pruning.
- **Ultimate Threshold:** The ultimate threshold dictates the absolute minimum correlation ratio value accepted as a match in the decomposition. To ensure an up-to-date reduced dictionary, it must be re pruned occasionally from the base dictionary. If the current match's correlation ratio violates the ultimate threshold, the match will be discarded, the dictionary will be re pruned, and the iteration will be repeated. Otherwise, the algorithm will extract the match and continue onto the next iteration. A violation of this threshold wastes computational time, but helps to maintain the integrity of the reconstruction.
- **Reprune Threshold:** This threshold is the second factor in the re pruning of the dictionary and helps to ensure smooth flow of the MPD++ algorithm. This threshold is typically set slightly above the ultimate threshold. As a result, the dictionary may be re pruned before the ultimate threshold is violated, thus conserving computational resources. If the current match's correlation ratio is below the re prune threshold and above the ultimate threshold, it will be subtracted from the signal and the dictionary will be re pruned at the beginning of the following iteration.

The interactions of the thresholds and the matched atom's correlation ratio determine: the action taken by the algorithm, the gains in computational performance, and the accuracy of the reconstruction. As long as the correlation ratio is above the ultimate threshold, the atom will be extracted. The optimal thresholds are dependent on the signal being decomposed. However, the thresholds may be approximated to produce significant increases in performance while maintaining the desired accuracy/performance emphasis. Figure 1 shows a flowchart for the CT technique. Dashed arrows and boxes indicate steps that differ from classical MPD.

Accuracy is often of primary importance. Unfortunately, standard MPD is often too computationally expensive to employ. Modifications that reduce the computational burden of the algorithm are desirable; however, reconstructions with an improved algorithm are often required be identical to those produced by classical MPD. When identical values are selected for the ultimate and reduction thresholds, the algorithm will discard any matches with a correlation ratio smaller than atoms previously rejected from the reduced dictionary. This prevents the extraction of any matches when an atom that was removed from the reduced dictionary may offer a better match.

If rapid decomposition is the primary focus, the ultimate threshold may be set at a value below the reduction threshold. This enables the improved algorithm to extract more atoms before having to re prune the base dictionary. However, it is possible for the algorithm to select an atom that normally would not have been picked or miss the true best match for that iteration because it was not present in the reduced dictionary. Significant gains in computational performance may be obtained at a decrease in accuracy that is acceptable for many applications.

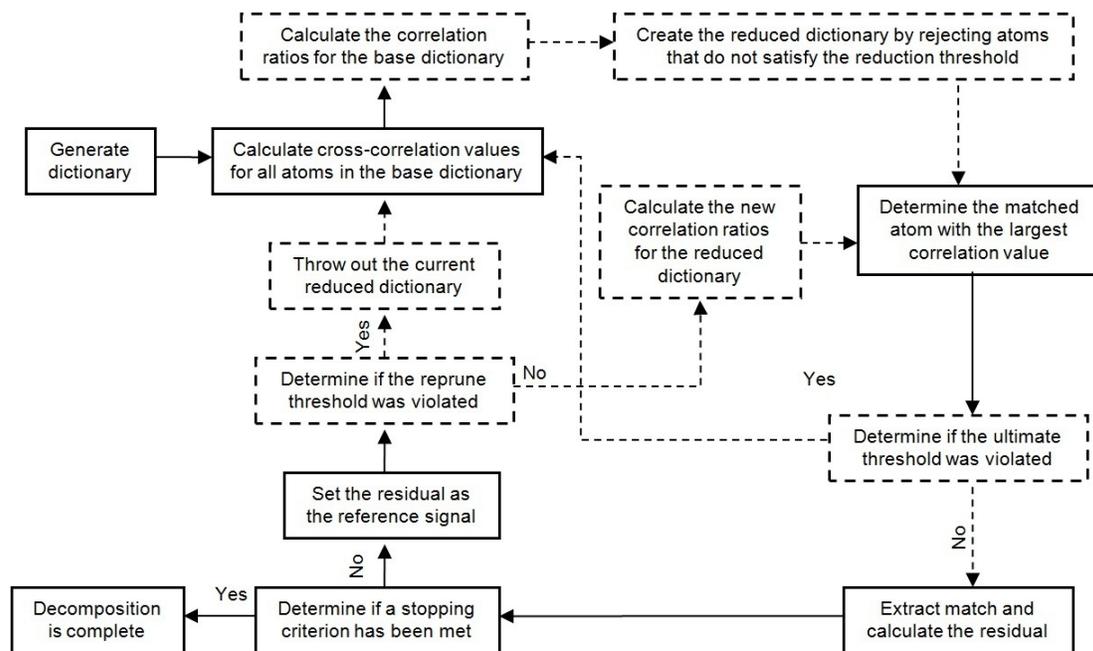


Figure 1: MPD++ Flowchart for Correlation Thresholding

Multiple Atom Extractions

With even the most effective, abridged dictionary, MPD is fundamentally limited by its iterative nature. The extraction of multiple atoms per iteration is proposed to reduce the computational burden of the algorithm. The Multiple Atom Extractions (MAE) modification alters the flow and organization of the MPD algorithm instead of merely addressing the encumbering effects caused by excessively large dictionaries.

Classical MPD determines and extracts the single best-correlated atom. However, MPD++ determines a specified number of atoms with the largest correlation values, referred to as the “top atoms”. Next, the algorithm subtracts the absolute best-correlated atom from the dictionary, referred to as the “primary atom”. All other top atoms are called “secondary atoms”. Then, the algorithm determines whether the next best correlated secondary atom overlaps the footprint of a previously removed match. If no overlap exists, the atom will also be subtracted from the signal. The process is repeated until all of the stored top atoms are either extracted or rejected. When this condition is reached, the algorithm will proceed onto the next iteration.

ATTEMPTED NUMBER OF ATOMS TO EXTRACT PER ITERATION

As the algorithm attempts to extract more atoms per iteration, the elapsed time and amount of extracted energy tend to decrease. For conservative, accurate reconstructions, extraction should only be attempted on a small number of atoms per iteration. For maximum performance, numerous atoms may be extracted in each iteration. In fact, if an excessively large number of top atoms are stored, the modification offers incredible performance increases with only a minor loss in reconstruction accuracy.

TOP ATOM TRACKING

Sometimes the algorithm will uncover atoms that overlap but model different parts of the signal, such as different frequency components. The “Top Atom Tracking” (TAT) mode offers previously rejected secondary atoms an additional chance to be extracted. If TAT is active, the algorithm will store an extra top atom to use as a reference; extraction will never be attempted on this extra stored atom. After the first pass through the list of the top atoms and the extraction of as many atoms as possible, the algorithm will recompute the correlation values for the top atoms that were rejected in the previous pass due to a footprint overlap. If these atoms have a correlation value larger than the extra top atom stored, the atoms will be extracted; otherwise these atoms will be permanently rejected from the current iteration.

RESULTS AND DISCUSSIONS

The results for the MPD++ algorithm with various combinations of the modifications are shown in Table I. The reconstruction accuracy is evaluated using normalized mean squared error (NMSE) as a metric. For the CT modification, the emphases of accuracy (“Acc”) and performance (“Perf”) are designated. The number of attempted atoms extractions is indicated by the column “AAE”. The reference and reconstructed signals are represented by x and \hat{x} respectively.

Table I: Observed reconstruction accuracies (NMSE) and run times of the MPD++ algorithm

Algorithm	Modifications				Accuracy $\sum_{i=1}^N \frac{(\hat{x}_i - x_i)^2}{N\sigma_x^2}$	Run time (sec)
	CT	MAE	AAE	TAT		
MPD	-	-	-	-	.0881	2078.4
MPD++	Acc	-	-	-	.0881	1092.3
MPD++	Perf	-	-	-	.0881	1044.2
MPD++	-	✓	5	-	.0882	1495.7
MPD++	-	✓	20	-	.0887	1162.2
MPD++	-	✓	50	-	.0887	970.9
MPD++	-	✓	140	-	.0889	808.1
MPD++	-	✓	5	✓	.0881	1216.0
MPD++	-	✓	20	✓	.0882	725.8
MPD++	-	✓	50	✓	.0884	454.0
MPD++	Acc	✓	5	✓	.0881	571.0
MPD++	Perf	✓	20	✓	.0882	348.9
MPD++	Perf	✓	140	✓	.0887	235.4

During the accurate and performance CT tests, time complexity reductions of 1.9 and 2 times, respectively, were observed. During the accuracy test, the integrity of the decomposition was guaranteed at a cost of performance; coincidentally, the performance mode was able to match the NMSE of accurate CT test.

The performance increases for the Multiple Atom Extractions technique depend on the number of attempted atom extractions per iteration. With five attempted atom extractions per iteration, a performance gain of 28 percent was observed. When decomposition was attempted on twenty and fifty atoms per iteration, performance gains of 44 and 53 percent were accompanied by subtle cost in reconstruction accuracy. As the extraction of more atoms per iteration is attempted, the performance gains tends to increase until an upper limit is reached. This boundary was determined, a test with 140 top atoms resulted in a 2.6 times reduction in the time complexity. Additional accuracy was sacrificed while producing no further performance gains when higher numbers were used. Top Atom Tracking significantly reduced the time complexity and appeared to slightly improve the accuracy of the reconstruction.

The MPD++ algorithm was tested with both modifications active. When these modifications emphasized accuracy, the algorithm extracted a similar amount of energy as classical MPD while requiring only 27.5% of the elapsed time. In the performance-centric tests, the MPD++ algorithm decomposed the signal 6 times faster than the standard algorithm. When the modified algorithm was pushed to its limits to determine the absolute maximum performance gains possible, the elapsed time was reduced by a factor of 8.8.

The tested signal consisted of 10,000 data points with features of interest stretching across a substantial portion of the signal. Due to its nature, the tested signal is particularly susceptible to experience dramatic performance gains with Multiple Atom Extractions and Correlation Thresholding. While most signals will experience significant performance gains, the degree of performance or accuracy enhancement is highly dependent on the signal.

CONCLUSION

The paper presents a variant of the MPD algorithm with dual emphases on accuracy and performance considered. The results confirm that the two modifications successfully improved the scalability and computational efficiency of the MPD algorithm. Correlation Thresholding decreased the time complexity by reducing the dictionary size. Multiple Atom Extractions also reduced the time complexity by decreasing the number of iterations required for a stopping criterion to be reached. Due to the nature of these modifications, they are capable of being stacked and have cumulative effects. The MPD++ algorithm was demonstrated using an over-complete dictionary on real life data. Reductions in computational time of 3.6 and 6 were observed for the emphases of accuracy and performance, respectively. Substantial performance increases, such as a 8.8 times improvement in time complexity, may be achieved at a slight cost in accuracy when knowledge of the signal is applied. The MPD++ algorithm was programmed using Matlab; the files are available on NASA's Discovery in Aeronautics System Health website (<https://dashlink.arc.nasa.gov>).

ACKNOWLEDGEMENTS

This research was conducted as part of an internship through NASA Undergraduate Student Research Program. The improved algorithm was tested on a signal courtesy of the Air Force Research Laboratory. This work was supported through funding from the NASA Aeronautics Research Mission Directorate, Aviation Safety Program, and Integrated Vehicle Health Management project. The authors thank Drs. Kanishka Bhaduri and Nikunj Oza for valuable discussions and suggestions.

BIBLIOGRAPHY

- [1] Das, S., I. Kyriakides, A. Chattopadhyay, and A. Papandreou-Suppappola. 2009. "Particle Filter Based Matching Pursuit Decomposition for Damage Quantification in Composite Structures," *Journal of Intelligent Material Systems and Structures*, 20(6):647–658.
- [2] Mallat, S. and Z. Zhang. 1993. "Matching Pursuits with Time-Frequency Dictionaries," *IEEE Transactions on Signal Processing*, 41(12):3397–3415.
- [3] Kristekova, M. and M. Kovacova. 2002. "New Version of Matching Pursuit Decomposition with Correct Representation of Linear Chirps," *Conference on Scientific Computing*, 33–41.
- [4] Imamura, K., Y. Koba, and H. Hashimoto. 2006. "A Fast Matching Pursuits Algorithm Using Sub-band Decomposition of Video Signals," *IEEE International Conference on Multimedia and Expo*, 729–732.
- [5] Schniter, P., L. C. Potter, and J. Ziniel. 2008. "Fast Bayesian Matching Pursuit: Model Uncertainty and Parameter Estimation for Sparse Linear Models," *IEEE Transactions on Signal Processing*.
- [6] Moschetti, F., L. Granai, P. Vandergheynst, and P. Frossard. 2002. "New Dictionary and Fast Atom Searching Method for Matching Pursuit Representation of Displaced Frame Difference," *Proceedings of the IEEE ICIP* 3:685–688.
- [7] Gribonval, R. 2001. "Fast Ridge Pursuit with a Multiscale Dictionary of Gaussian Chirps," *IEEE Transactions on Signal Processing*, 49(5):994–1001.